arXiv:1411.1285v1 [stat.ML] 5 Nov 2014

Controlling false discoveries in high-dimensional situations

Boosting with stability selection

Benjamin Hofner*† Luigi Boccuto † Markus Göker‡

November 6, 2014

Abstract

Modern biotechnologies often result in high-dimensional data sets with much more variables than observations ($n \ll p$). These data sets pose new challenges to statistical analysis: Variable selection becomes one of the most important tasks in this setting. We assess the recently proposed flexible framework for variable selection called stability selection. By the use of resampling procedures, stability selection adds a finite sample error control to high-dimensional variable selection procedures such as Lasso or boosting. We consider the combination of boosting and stability selection and present results from a detailed simulation study that provides insights into the usefulness of this combination. Limitations are discussed and guidance on the specification and tuning of stability selection is given. The interpretation of the used error bounds is elaborated and insights for practical data analysis are given. The results will be used to detect differentially expressed phenotype measurements in patients with autism spectrum disorders. All methods are implemented in the freely available R package **stabs**.

KEYWORDS boosting, error control, variable selection, stability selection

1. Introduction

Variable selection is a notorious problem in many applications. The researcher collects many variables on each study subject and then wants to identify the variables that have an influence on the outcome variable. This problem becomes especially pronounced with modern high-throughput experiments where the number of variables p is often much larger than the number of observations n (e.g., genomics, transcriptomics, proteomics, metabolomics, metabonomics and phenomics; see Chaturvedi, Goeman, Boer, van Wieringen, and de Menezes 2014; Wang, Gerstein, and Snyder 2009; Mallick and Kuster 2010; Ludwig and Günther 2011; Lindon, Holmes, and Nicholson 2003; Groth, Weiss, Pohlenz, and Leser 2008). One of the major aims in the analysis of these high-dimensional

*E-mail: benjamin.hofner@fau.de
†Department of Medical Informatics, Biometry and Epidemiology, Friedrich-Alexander-Universität Erlangen-Nürnberg, Waldstraße 6, 91054 Erlangen, Germany
†Greenwood Genetic Center, 113 Gregor Mendel Circle, Greenwood, SC 29646, USA
‡Leibniz Institute DSMZ – German Collection of Microorganisms and Cell Cultures, Inhoffenstraße 7b, 38124 Braunschweig, Germany

data sets is to detect the signal variables S, while controlling the number of selected noise variables N. Stepwise regression models are a standard approach to variable selection in settings with relatively few variables. However, even in this case this approach is known to be very unstable (see e.g., Flack and Chang 1987; Austin and Tu 2004; Austin 2008). Recent approaches that try to overcome this problem and can also be used in high-dimensional settings with $n \ll p$ include penalized regression approaches such as the lasso (Tibshirani 1996; Efron, Hastie, Johnstone, and Tibshirani 2004), elastic net (Zou and Hastie 2005), and boosting (Friedman, Hastie, and Tibshirani 2000), or tree based approaches such as random forests (Breiman 2001; Strobl, Boulesteix, Zeileis, and Hothorn 2007). More recently, Meinshausen and Bühlmann (2010) proposed stability selection, an approach based on resampling of the data set which can be combined with many selection procedures and is especially useful in high-dimensional settings. Stability selection has since been widely used, e.g. for gene regulatory network analysis (Haury, Mordelet, Vera-Licona, and Vert 2012; Marbach, Costello, Küffner, Vega, Prill, Camacho, Allison, Kellis, Collins, Stolovitzky *et al.* 2012), in genome-wide association studies (He and Lin 2011), graphical models (Fellinghauer, Bühlmann, Ryffel, von Rhein, and Reinhardt 2013; Bühlmann, Kalisch, and Meier 2014) or even in ecology (Hothorn, Müller, Schröder, Kneib, and Brandl 2011). In most publications, stability selection is used in combination with lasso or similar penalization approaches. Here, we discuss the combination of stability selection with component-wise functional gradient descent boosting (Bühlmann and Yu 2003) which allows one to specify competing effects, which are subject to selection, more flexibly. For details on functional gradient descent boosting, see Bühlmann and Hothorn (2007) and Hofner, Mayr, Robinzonov, and Schmid (2014b).

We will provide a short, rather non-technical introduction to boosting in Section 2. Stability selection, which controls the per-family error rate, will be introduced in Section 3, where we also give an overview on common error rates and some guidance on the choice of the parameters in stability selection. Section 4 presents an empirical evaluation of boosting with stability selection. In our case study (Section 5) we will examine autism spectrum disorder (ASD) patients and compare them to healthy controls using the boosting approach in conjunction with stability selection. The aim is to detect differentially expressed phenotype measurements. More specifically, we try to assess which amino acid pathways differ between healthy subjects and ASD patients.

2. A Short Introduction to Boosting

Consider a generalized linear model

$$\mathbb{E}(y|\mathbf{x}) = h(\eta(\mathbf{x})) \tag{1}$$

with outcome y, appropriate response function h and linear predictor $\eta(x)$. Let the latter be defined as

$$\eta(\mathbf{x}) = \beta_0 + \sum_{j=1}^{p} \beta_j x_j, \tag{2}$$

with covariates $\mathbf{x} = (x_1, \ldots, x_p)$, and corresponding effects $\beta_j,\ j = 0, \ldots, p$. Model fitting aims at minimizing the expected loss $\mathbb{E}(\rho(y, \mathbf{x}))$ with an appropriate loss function $\rho(y, \mathbf{x})$. The loss function is defined by the fitting problem at hand. Thus, for example, Gaussian regression models, i.e. least squares regression models, aim to minimize the squared loss $\rho(y, \mathbf{x}) = (y - \eta(\mathbf{x}))^2$. Generalized

linear models can be obtained by maximizing the log-likelihood or, analogously, by minimizing the negative log-likelihood function. Logistic regression models with binary outcome, for example, can be obtained by using the negative binomial log-likelihood

$$\rho(y, \mathbf{x}) = -y \log(P(y = 1|\mathbf{x})) + (1 - y) \log(1 - P(y = 1|\mathbf{x}))$$

as loss function or a reparametrization thereof (Bühlmann and Hothorn 2007).

In practice, one cannot minimize the expected loss function. Instead, we optimize the empirical risk function

$$\mathcal{R}(\mathbf{y}, \mathbf{X}) = n^{-1} \sum_{i=1}^{n} \rho(y_i, \eta(\mathbf{x}_i)) \tag{3}$$

with observations $\mathbf{y} = (y_1, \ldots, y_n)^\top$ and $\mathbf{X} = (\mathbf{x}_1^\top, \ldots, \mathbf{x}_n^\top)^\top$. This can be done for arbitrary loss functions by component-wise functional gradient descent boosting (Bühlmann and Yu 2003). The algorithm is especially attractive owing to its intrinsic variable selection properties (Kneib, Hothorn, and Tutz 2009; Hofner, Hothorn, Kneib, and Schmid 2011).

One begins with a constant model $\hat{\eta}^{[0]}(\mathbf{x}_i) \equiv 0$ and computes the residuals $\mathbf{u}^{[1]} = (u_1^{[1]}, \ldots, u_n^{[1]})^\top$ defined by the negative gradient of the loss function

$$u_i^{[m]} := - \left. \frac{\partial \rho(y_i, \eta)}{\partial \eta} \right|_{\eta = \hat{\eta}^{[m-1]}(\mathbf{x}_i)} \tag{4}$$

evaluated at the fit of the previous iteration $\hat{\eta}^{[m-1]}(\mathbf{x}_i)$ (see Bühlmann and Yu 2003; Bühlmann and Hothorn 2007; Hothorn, Bühlmann, Kneib, Schmid, and Hofner 2010). Each variable $x_1, \ldots, x_p$ is fitted separately to the residuals $\mathbf{u}^{[m]}$ by least squares estimation (this is called the "base-learner"), and only the variable j^* that describes these residuals best is updated by adding a small percentage ν of the fit $\hat{\beta}_{j^*}$ (e.g., $\nu = 10\%$) to the current model fit, i.e.,

$$\hat{\eta}^{[m]} = \hat{\eta}^{[m-1]} + \nu \cdot \hat{\beta}_{j^*}.$$

New residuals $\mathbf{u}^{[m+1]}$ are computed, and the whole procedure is iterated until a fixed number of iterations $m = m_{\text{stop}}$ is reached. The final model $\hat{\eta}^{[m_{\text{stop}}]}(\mathbf{x}_i)$ is defined as the sum of all models fitted in this process. Instead of using linear base-learners (i.e., linear effects) to fit the negative gradient vector $\mathbf{u}^{[m]}$ in each boosting step, one can also specify smooth base-learners for the variables x_j (see e.g. Schmid and Hothorn 2008), which are then fitted by penalized least squares estimation. As we update only one modeling component in each boosting iteration, variables are selected by stopping the boosting procedure after an appropriate number of iterations ("early stopping"). This number is usually determined using cross-validation techniques (see e.g., Mayr, Hofner, and Schmid 2012).

3. Stability Selection

A problem of many statistical learning approaches including boosting with early stopping is that despite regularization one often ends up with relatively rich models (Mayr *et al.* 2012; Meinshausen and Bühlmann 2010). A lot of noise variables might be erroneously selected. To improve the selection process and to obtain an error control for the number of falsely selected noise variables

Meinshausen and Bühlmann (2010) proposed stability selection. This is a versatile approach, which can be combined with all high-dimensional variable selection approaches. Stability selection is based on sub-sampling and controls the *per-family error rate* $\mathbb{E}(V)$, where V is the number of false positive variables (for more details on error rates see Section 3.2).

Consider a data set with p predictor variables x_j, $j = 1, \ldots, p$ and an outcome variable y. Let $S \subseteq \{1, \ldots, p\}$ be the set of signal variables, and let $N \subseteq \{1, \ldots, p\}/S$ be the set of noise variables. The set of variables that are selected by the statistical learning procedure is denoted by $\hat{S}_n \subseteq \{1, \ldots, p\}$. This set $\hat{S}_n$ can be considered to be an estimator of S, based on a data set with n observations. In short, for stability selection with boosting one proceeds as follows:

1.) Select a random subset of size $\lfloor n/2 \rfloor$ of the data, where $\lfloor x \rfloor$ denotes the largest integer $\leq x$.

2.) Fit a boosting model and continue to increase the number of boosting iterations m_{stop} until q base-learners are selected. $\hat{S}_{\lfloor n/2 \rfloor, b}$ denotes the set of selected variables.

3.) Repeat the steps 1) and 2) for $b = 1, \ldots, B$.

4.) Compute the relative selection frequencies

$$\hat{\pi}_j := \frac{1}{B} \sum_{b=1}^{B} \mathbb{I}_{\{j \in \hat{S}_{\lfloor n/2 \rfloor, b}\}} \tag{5}$$

per variable (or actually per base-learner).

5.) Select all base-learners that were selected with a frequency of at least π_{thr}, where π_{thr} is a pre-specified threshold value. Thus, we obtain a set of *stable variables* $\hat{\text{S}}_{\text{stable}} := \{j : \hat{\pi}_j \geq \pi_{\text{thr}}\}$.

Meinshausen and Bühlmann (2010) show that this selection procedure controls the per-family error rate (*PFER*). An upper bound is given by

$$\mathbb{E}(V) \leq \frac{q^2}{(2\pi_{\text{thr}} - 1)p} \tag{6}$$

where q is the number of selected variables per boosting run, p is the number of (possible) predictors and π_{thr} is the threshold for selection probability. The theory requires two assumptions to ensure that the error bound holds:

(i) The distribution $\{\mathbb{I}_{\{j \in \hat{\text{S}}_{\text{stable}}\}}, j \in N\}$ needs to be exchangeable for all noise variables N.

(ii) The original selection procedure, boosting in our case, must not be worse than random guessing.

In practice, assumption (i) essentially means that each noise variable has the same selection probability. Thus, all *noise variables* should, for example, have the same correlation with the signal variables (and the outcome). For examples of situations where exchangeability is given see Meinshausen and Bühlmann (2010). Assumption (ii) means that signal variables should be selected with higher probability than noise variables. This assumption is usually not very restrictive as we would expect it to hold for any sensible selection procedure.

Choice of parameters The stability selection procedure mainly depends on two parameters: the number of selected variables per boosting model q and the threshold value for stable variables π_{thr}. Meinshausen and Bühlmann (2010) propose to chose $\pi_{\text{thr}} \in (0.6, 0.9)$ and claim that the threshold has little influence on the selection procedure. In general, any value $\in (0.5, 1)$ is potentially acceptable, i.e. a variable should be selected in more than half of the fitted models in order to be considered stable. The number of selected variables q should be chosen so high that in theory all signal variables S can be chosen. If q was too small, one would inevitably select only a small subset of the signal variables S in the set $\hat{\text{S}}_{\text{stable}}$ as $|\hat{\text{S}}_{\text{stable}}| \leq |\hat{S}_{\lfloor n/2 \rfloor, b}| = q$ (if $\pi_{\text{thr}} > 0.5$).

The choice of the number of subsamples B is of minor importance as long as it is large enough. Meinshausen and Bühlmann (2010) propose to use $B = 100$ replicates, which seems to be sufficient for an accurate estimation of $\hat{\pi}_j$ in most situations.

In general, we would recommend to choose an upper bound $PFER_{\text{max}}$ for the $PFER$ and specify either q or π_{thr}, preferably π_{thr}. The missing parameter can then be computed from Equation (6), where equality is assumed. In a second step, one should check that the computed value is sensible, i.e. that $\pi_{\text{thr}} \in (0.5, 1)$, or that q is not too small, or that $PFER_{\text{max}}$ is not too small or too large. Note that the $PFER$ can be greater than one as it resembles the tolerable expected number of falsely selected noise variables. An overview on common error rates is given in Section 3.2, where we also give some guidance on the choice of $PFER_{\text{max}}$.

The size of the subsamples is no tuning parameter but should always be chosen to be $\lfloor n/2 \rfloor$. This an essential requirement for the derivation of the error bound (6) as can be seen in the proof of Lemma 2 (Meinshausen and Bühlmann 2010), which is used to proof the error bound. Other (larger) subsample sizes would theoretically be possible but would require the derivation of a different error bound for that situation.

3.1. Improved Version of Stability Selection

A modification of stability selection was introduced by Shah and Samworth (2013). One major difference to the original stability selection approach is that instead of using B independent subsamples of the data, Shah and Samworth use $2B$ complementary pairs: One draws B subsamples of size $\lfloor n/2 \rfloor$ from the data and uses, for each subsample, the remaining observations as a second complementary subsample.

More importantly, error bounds are theoretically derived that hold without assuming exchangeability of the noise variables (and without assuming that the original selection procedure is not worse than random guessing). The drawback of being able to drop these assumptions (i) and (ii)) is that the modified bounds do not control the per-family error rate, but the *expected number of selected variables with low selection probability*

$$\mathbb{E}(|\hat{\text{S}}_{\text{stable}} \cap L_\theta|), \tag{7}$$

where $\hat{\text{S}}_{\text{stable}}$ denotes the set of variables selected by stability selection, and $L_\theta = \{j : \hat{\pi}_j \leq \theta\}$ denotes the set of variables that have a low selection probability under $\hat{S}_{\lfloor n/2 \rfloor}$, i.e. a selection probability below θ in one boosting run on a subsample of size $\lfloor n/2 \rfloor$. Usually, this threshold for low selection probabilities is chosen as $\theta = \frac{q}{p}$, i.e. the average fraction of selected variables. Thus, this error rate represents the expected number of variables that are unlikely to be selected but are selected.

Here, the selection probability $\hat{\pi}_j$ (Eq. 5) needs to be computed over all $2B$ random (complemen-

tary) subsamples. Additionally, let the simultaneous selection probability $\widetilde{\pi}_j$ be defined as follows (Shah and Samworth 2013):

$$\widetilde{\pi}_j := \frac{1}{B} \sum_{b=1}^{B} \mathbb{I}_{\{j \in \hat{S}_b^1\}} \cdot \mathbb{I}_{\{j \in \hat{S}_b^2\}}, \tag{8}$$

where $\mathbb{I}_{\{j \in S\}}$ is the indicator function which is one if $j \in S$ and zero otherwise. $\hat{S}_b^1$ is the set of selected variables on the bth random subset of size $\lfloor n/2 \rfloor$ and $\hat{S}_b^2$ is the selection on the complementary pair of this random subset. Note that both sets of selected variables are derived with the original learning procedure without applying the stability selection threshold so far.

Shah and Samworth (2013) derive three error bounds for the *expected number of low selection probability variables*:

(E1) A worst case error bound is derived for all $\pi_{\text{thr}} \in (0.5, 1]$:

$$\mathbb{E}(|\,\hat{\mathrm{S}}_{\text{stable}} \cap L_\theta|) \leq \frac{\theta}{2\pi_{\text{thr}} - 1}\, \mathbb{E}(|\hat{S}_{\lfloor n/2 \rfloor} \cap L_\theta|) \leq \frac{\theta}{2\pi_{\text{thr}} - 1}\, q \tag{9}$$

If $\theta = \frac{q}{p}$, this error bound is equal to (6) but does not require that assumptions (i) and (ii) hold.

(E2) A second, tighter, error bound assumes that the simultaneous selection probabilities $\widetilde{\pi}_j$ have a unimodal probability distribution for all $j \in L_\theta$. If additionally $\theta \leq 1/\sqrt{3} \approx 0.577$ holds, the error bound can be written as

$$\mathbb{E}(|\,\hat{\mathrm{S}}_{\text{stable}} \cap L_\theta|) \leq \frac{\theta}{c(\pi_{\text{thr}}, B)}\, \mathbb{E}(|\hat{S}_{\lfloor n/2 \rfloor} \cap L_\theta|) \leq \frac{\theta}{c(\pi_{\text{thr}}, B)}\, q \tag{10}$$

with constant

$$c(\pi_{\text{thr}}, B) = \begin{cases} 2\left(2\pi_{\text{thr}} - 1 - \frac{1}{2B}\right) & \text{if } \pi_{\text{thr}} \in (c_{\min}, \frac{3}{4}] \\ \frac{1+1/B}{4\left(1-\pi_{\text{thr}}+\frac{1}{2B}\right)} & \text{if } \pi_{\text{thr}} \in (\frac{3}{4}, 1], \end{cases}$$

and $c_{\min} = \min(\frac{1}{2} + \theta^2, \frac{1}{2} + \frac{1}{2B} + \frac{3}{4}\theta^2)$. One needs to further assume that $\pi_{\text{thr}} \in \left\{\frac{1}{2} + \frac{2}{2B}, \frac{1}{2} + \frac{3}{2B}, \ldots, 1\right\}$ for the bound to hold. However, this is no restriction in practice, as for typical values of B such as $B = 50$ or $B = 100$, all values of $\pi_{\text{thr}} \geq 0.51$ in steps of 0.01 or $\pi_{\text{thr}} \geq 0.505$ in steps of 0.005, respectively, are permitted.

(E3) The third error bound assumes that the simultaneous selection probabilities $\widetilde{\pi}_j$ have an r-concave probability distribution with $r = -\frac{1}{2}$ and that the selection probabilities $\hat{\pi}_j$ have an r-concave probability distribution with $r = -\frac{1}{4}$ for all $j \in L_\theta$. With f_j being the distribution of $\widetilde{\pi}_j$ and g_j being the distribution of $\hat{\pi}_j$ this is equivalent to the assumptions that $f_j^{-1/2}$ and $g_j^{-1/4}$ must be convex. The r-concavity assumption lies in between unimodality and the stronger log-concavity assumption. For details on r-concavity we refer to Shah and Samworth (2013). If the r-concavity assumption holds, the error bound can be further refined as

$$\begin{aligned} \mathbb{E}(|\,\hat{\mathrm{S}}_{\text{stable}} \cap L_\theta|) &\leq \min\left\{D\left(2\pi_{\text{thr}} - 1; \theta^2, B, -\frac{1}{2}\right), D\left(\pi_{\text{thr}}; \theta, 2B, -\frac{1}{4}\right)\right\} |L_\theta| \\ &\leq \min\left\{D\left(2\pi_{\text{thr}} - 1; \theta^2, B, -\frac{1}{2}\right), D\left(\pi_{\text{thr}}; \theta, 2B, -\frac{1}{4}\right)\right\} p. \end{aligned} \tag{11}$$

The function $D(\xi; \theta, B, r)$ denotes the maximum of the probability $P(X \leq \xi)$ with $\mathbb{E}(X) \leq \theta$

over all r-concave random variables X on a discrete support $\{0, 1/B, 2/B, \ldots, 1\}$. For details see Shah and Samworth (2013, Appendix A.4).

With these additional assumptions we get much tighter error bounds. The reason for tighter bounds can be found in the application of refined bounds in Markov's inequality that make use of the distributional assumptions. Markov's inequality is used on the simultaneous selection probabilities $\widetilde{\pi}_j$ in the derivation of the error bounds (see Shah and Samworth 2013, App. A.1–A.3).

One should be aware that the assumptions are on the *distribution* of the selection probabilities and not on the selection probability itself. The unimodality assumption seems to generally hold in practice. The r-concavity assumption may fail, if the number of subsamples B increases, since as B increases, r-concavity requires an increasing number of inequalities to hold for the distribution of $\widetilde{\pi}_j$. However, the same problem does not occur for the unimodal bound, and when $B = 50$, the bounds constructed using the r-concavity assumption seem to hold in a wide variety of scenarios (Shah, 2014, personal communication; see also Section 4.1).

Interpretation of Error Bounds (E1) – (E3) If the exchangeability assumption holds and the selection procedure is not worse than random guessing, then all noise variables have a "below average" selection probability. Hence, the low selection probability variables will include all noise variables, i.e. $L_\theta = N$. Controlling the *expected number of selected variables with low selection probability* is thus in this case identical to controlling the expected number of false positives:

$$\mathbb{E}(|\,\hat{\mathbf{S}}_{\text{stable}} \cap L_\theta|) = \mathbb{E}(|\,\hat{\mathbf{S}}_{\text{stable}} \cap N|) = \mathbb{E}(V).$$

Stability selection can consequently be thought to control the per-family error rate in all three cases (E1) – (E3). On the other hand, if exchangeability does not hold, this means that we have "special" noise variables, e.g., noise variables that are stronger correlated with signal variables than other noise variables. If this correlation is so strong that a variable is selected with "above average selection probability", it is difficult to think of this variable as noise variables anyway. Thus controlling the *expected number of selected variables with low selection probability* is again similar or even practically identical to controlling the expected number of false positives.

3.2. Definitions and Discussion of Common Error Rates

There are various definitions of error rates that are used in statistics, especially in the case of multiple testing. Let m be the number of tested hypothesis, R the number of rejected hypothesis and V the number falsely rejected hypotheses as defined above (cf. Benjamini and Hochberg 1995). In our case, m is the number of predictor variables p or more general the number of base-learners in the boosting model. Commonly used error rates include the per-comparison error rate $PCER = \mathbb{E}(V)/m$, the per-family error rate $PFER = \mathbb{E}(V)$, the family-wise error rate $FWER = \mathbb{P}(V \geq 1)$, and the false discovery rate $FDR = \mathbb{E}(\frac{V}{R})$ (Benjamini and Hochberg 1995). The per-comparison error rate is the standard error rate without adjustment for multiplicity.

For a given test situation it holds that

$$PCER \leq FWER \leq PFER\,.$$

Thus, for a fixed significance level α it holds that *PFER*-control is more conservative than *FWER*-control which is in turn more conservative than *PCER*-control (Dudoit, Shaffer, and Boldrick 2003). The *FDR*, which is often used in (very) high-dimensional settings such as gene expression studies uses another error definition by relating the number of false discoveries to the number of rejected null hypotheses. One can show that in a given test situation

$$FDR \leq FWER,$$

and thus for a fixed level α, *FWER*-control is more conservative than *FDR*-control (Dudoit *et al.* 2003). In conclusion, it holds that $FDR \leq FWER \leq PFER$. Controlling the *PFER* is a (very) conservative approach for controlling errors in multiple testing situations. Hence, a procedure that controls the *PFER* at a certain level α also controls all other error rates discussed in this section at this level. Obviously the error bound will be very conservative upper bound for both the *FWER* and *FDR*.

The standard approach for hypotheses testing, neglecting multiplicity, would be to specify a bound for the per-comparison error rate by using a significance level α, e.g. $\alpha = 0.05$. This is equal to specifying $PFER_{\max} \leq m\alpha$. This provides some guidance on how to choose an upper bound for the *PFER*: Usually, $\alpha \leq PFER_{\max} \leq m\alpha$ seems a good choice, where $PFER_{\max} = \alpha$ would (conservatively) control the *FWER* on the level α, while $PFER_{\max} = m\alpha$ would control the unadjusted per-comparison error rate on the level α. Everything in between can be considered to control the *PCER* on the level α "with some multiplicity adjustment".

4. Empirical Evaluation

To evaluate the impact of the tuning parameters q and π_{thr}, the upper bound $PFER_{\max}$, and the assumptions for the computation of the upper bound on the selection properties, we conducted a simulation study using boosting in conjunction with stability selection. Additionally, we examined the impact of the characteristics of the data set on the performance.

We considered a classification problem with a binary outcome variable. The data were generated according to a linear logistic regression model with linear predictor $\eta = \mathbf{X}\beta$ and

$$Y \sim \text{Binom}\left(\frac{\exp(\eta)}{1+\exp(\eta)}\right).$$

The observations $x_i = (x_{i1}, \ldots, x_{ip})$, $i = 1, \ldots, n$ were independently drawn from

$$x \sim \mathcal{N}(0, \Sigma),$$

and gathered in the design matrix $\mathbf{X}$. We set the number of predictor variables to $p \in \{100, 500, 1000\}$, and the number of observations to $n \in \{50, 100, 500\}$. The number of influential variables varied within $p_{\text{infl}} \in \{2, 3, 8\}$, where β_j was sampled from $\{-1, 1\}$ for an influential variable and set to zero for all non-influential variables. We used two settings for the design matrix:

(a) independent predictor variables, i.e. $\Sigma = \mathbf{I}$,

(b) and correlated predictor variables drawn from a Toeplitz design with covariance matrix $\Sigma_{kl} = 0.9^{|k-l|}$, $k, l = 1, \ldots, p$.

For each of the data settings we used all three error bounds in combination with varying parameters $\pi_{\text{thr}} \in \{0.6, 0.75, 0.9\}$, and $PFER_{\max} \in \{0.05, 1, 2, 5\}$. We used the standard subsampling scheme with $B = 100$ subsamples for the error bound (E1) and complementary pairs with $B = 50$ subsamples for the improved error bounds (E2) and (E3). Each simulation setting was repeated 50 times.

4.1. Results

Figure 1 displays the true positive rates for different $PFER_{\max}$ bounds, the three assumptions (E1) to (E3) and for the two correlation schemes. Different sizes of the data set (n and p) as well as different numbers of true positives (p_{infl}) were not depicted as separate boxplots. For each upper bound $PFER_{\max}$ and each data situation (uncorrelated/Toeplitz), the true positive rate (TPR) increased with stronger assumptions (E1) to (E3). The true positive rate was lower when the predictors were correlated.

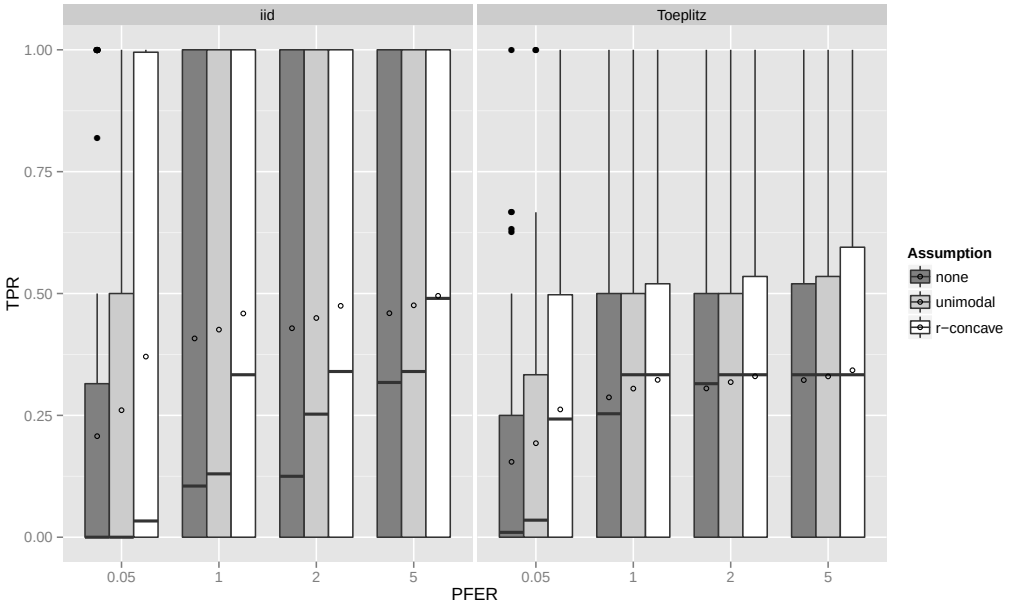

Figure 1: Boxplots for the true positives rates (TPR) for all simulation settings with separate boxplots for the correlation settings (independent predictor variables or Toeplitz design), $PFER_{\max}$ and the assumption used to compute the error bound. Each observation in the boxplot is the average of the 50 simulation replicates. The open circles in the boxes represent the average true positive rates.

If the number of observations n increased, the TPR increased as well with more extreme cases for uncorrelated predictors (see Appendix; Figure 5). With very few observations ($n = 50$), the TPR was generally very small. Considering the size of the subsamples, which is equal to 25, this is quite natural. Recently, Schmid, Hothorn, Krause, and Rabe (2012) advocated to increase the sample size of the subsamples from $\lfloor n/2 \rfloor$ to larger values to avoid biased selection of base-learners due to too small samples. Yet, as discussed above, this is currently not possible, as one would need to derive a different error bound for that situation. Conversely, the TPR decreases with an increasing number of truly influential variables p_{infl} (see Appendix; Figure 6). The threshold π_{thr} is less important (see Appendix; Figure 7), as long as it is large enough to result in enough variables q to be selected and

not too large so that too many variables would be selected in each run (see Appendix; Figure 8). Note that the dependence on the threshold is stronger in the case with correlated observations.

The number of false positives, which is bounded by the upper bound for the per-family error rate, is depicted in Figure 2. Overall, the error rate seemed to be well controlled with some violations of the less conservative bounds in the median settings. However, overall the error bound was violated in only 1.2 % (4 cases) under the unimodality assumption and 4.0 % (13 cases) under the r-concavity assumption. Especially the standard error bound (E1) seemed to be conservatively controlled. The average number of false positives increased with increasing $PFER_{\text{max}}$ and with stronger distributional assumptions on the simultaneous selection probabilities. In general, one should note that stability selection is quite conservative as it controls the *PFER*. The given upper bounds for the *PFER* corresponded to per-comparison error rates between 0.05 and 0.00005.

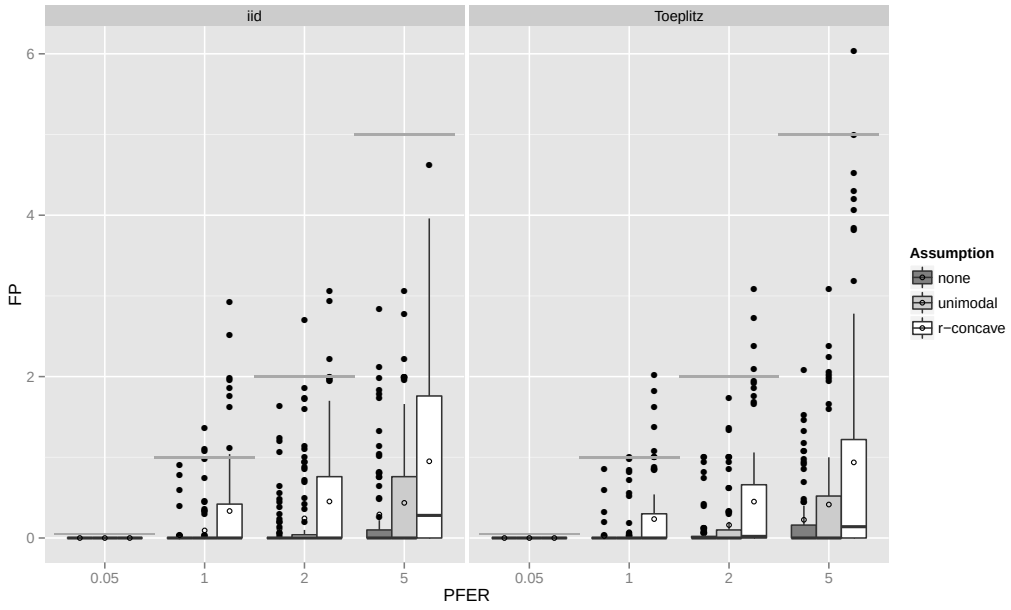

Figure 2: Boxplots for the number of false positives for all simulation settings with separate boxplots for the correlation settings (independent predictor variables or Toeplitz design), $PFER_{\text{max}}$ and the assumption used to compute the error bound. Each observation in the boxplot is the average of the 50 simulation replicates. The open circles represent the average number of false positives. The grey horizontal lines represent the error bounds.

If the number of observations n increased, the number of false positives decreased on average but the variability increased as well (see Appendix; Figure 9). The number of false positives showed a tendency to decrease with an increasing number of truly influential variables p_{infl} (see Appendix; Figure 10). If the threshold π_{thr} was larger, i.e., only highly frequently selected variables were considered to be stable, the number of false positives decreased (see Appendix; Figure 11). Yet, considering the corresponding number of selected variables per boosting run q (which is inversely related to the threshold π_{thr}), one could see that not only large values of q lead to low numbers of false positives but also small values (12). This observation is somehow contrary to the optimal choices of q with respect to the true positive rate. However, an optimal true positive rate is more important than a low number of false positives as long as the error rate is controlled.

5. Case Study: Differential Phenotype Expression for ASD patients versus controls

We examined autism spectrum disorder (ASD) patients (Manning-Courtney, Murray, Currans, Johnson, Bing, Kroeger-Geoppinger, Sorensen, Bass, Reinhold, Johnson, and Messerschmidt 2013) and compared them to healthy controls. The aim was to detect differentially expressed amino acid pathways, i.e. amino acid pathways that differ between healthy subjects and ASD patients (Boccuto, Chen, Pittman, Skinner, McCartney, Jones, Bochner, Stevenson, and Schwartz 2013). We used measurements of absorbance readings from Phenotype Microarrays developed by Biolog (Hayward, CA). The arrays are designed so as to expose the cells to a single carbon energy source per well and evaluate the ability of the cells to utilize this energy source to generate NADH (Bochner, Gadzinski, and Panomitros 2001). The array plates were incubated for 48 h at 37°C in 5% CO2 with 20,000 lymphoblastoid cells per well. After this first incubation, Biolog Redox Dye Mix MB was added (10 μL/well) and the plates were incubated under the same conditions for an additional 24 h. As the cells metabolize the carbon source, tetrazolium dye in the media is reduced, producing a purple color according to the amount of NADH generated. At the end of the 24 h incubation, the plates were analyzed utilizing a microplate reader with readings at 590 and 750 nm. The first value (A_{590}) indicated the highest absorbance peak of the redox dye and the second value (A_{750}) gave a measure of the background noise. The relative absorbance ($A_{590-750}$) was calculated per well.

Each row of the data set described the measurement of *one well per biological replicate*. With $n = 35$ biological replicates (17 ASD patients and 18 controls) and $p = 4 \cdot 96 = 384$ wells we thus theoretically got $n \cdot p = 13440$ observations. Due to one missing value the data set finally contained only 13439 observations. The data is available as a supplement to Boccuto *et al.* (2013) and in the R package **opm** (Vaas, Sikorski, Hofner, Buddruhs, Fiebig, Klenk, and Göker 2013, Göker 2014), which was also used to store, manage and annotate the data set.

For all available biological replicates we obtained the amino acid annotation for each measurement in that replicate, i.e. we set up an incidence vector per observation for all available peptides. The incidence vector was one if the peptide contained that amino acid and zero if it did not. We ended up with 27 amino acid occurrence annotations in total (including some non-proteinogenic amino acids). In the next step, we modeled the differences of the measured values between ASD patients and controls to assess which amino acid pathways were differentially expressed. Therefore we set up a model of the following form:

$$\begin{aligned}\log(y) = \beta_0 + \beta_1 \text{group} + b_{\text{id}} + \beta_{2,1} I_{\text{P1}} + \beta_{2,2} I_{\text{P2}} + \ldots + \\ + X(\text{group}) \cdot \widetilde{b}_{\text{id}} + X(\text{group}) \cdot \beta_{3,1} I_{\text{P1}} + X(\text{group}) \cdot \beta_{3,2} I_{\text{P2}} + \ldots,\end{aligned}$$

where y was the measured PM value, β_0 was an overall intercept, β_1 was the overall group effect (the difference between ASD patients and controls irrespective of the amino acid that the measurement belonged to). Additionally, we used an random effect for the replicate (b_{ID}) to account for subject-specific effects. The amino acid effects $\beta_{2,j}$ represent the differences of the $\log(y)$ values between amino acid, as $I_{\text{P}j}$ is an indicator function, which was 0 if the well did not belong to amino acid j, and 1 if it did; this means we obtained dummy-coded effect estimates from the first line of the model formula.

The most interesting part was given by the second line of the model: $X(\text{group})$ was a group-specific function which was either -1 for controls or 1 for ASD cases. We used this sum-to-zero constraint in an interaction with dummy-coded amino acid effects. The coefficients $\beta_{3,j}$ hence represented the deviation of the groups from the global effect of the jth amino acid. If $\beta_{3,j} = 0$, no group-specific effect was present, i.e. the amino acid did not differ between the groups. If $\beta_{3,j} \neq 0$, the difference between the two groups was twice this effect, i.e. $X(\text{ASD}) \cdot \beta_{3,j} - (X(\text{Control}) \cdot \beta_{3,j}) = 1 \cdot \beta_{3,j} - (-1 \cdot \beta_{3,j}) = 2\beta_{3,j}$. Note that we also specified a group-specific random effect $\widetilde{b}_{\text{ID}}$.

First, we fitted an offset model containing all main effects, i.e. we modeled differences in the maximum curve height with respect to different amino acids while neglecting possible differences in amino acid effects between groups. In a second step, we started from this offset model and additionally allowed for interactions between the group and the amino acids, while keeping the main effects in the list of possible base-learners, and checked if any interactions were present. These represent differential PM expressions between groups.

In total, we ended up with 57 base-learners (group effect, main amino acid effects, group-specific effects, and an overall and a group-specific random effect). All models were fitted using boosting. The selection of differentially expressed amino acids was done using stability selection. We set the number of selected variables per boosting model to $q = 10$ and chose an upper bound for the $PFER \leq 1$. To judge the magnitude of the multiplicity correction, we related the used $PFER$ to the significance level α, i.e. the standard $PCER$: The upper bound for the $PFER$ equaled $\alpha = 1/57 = 0.0175$ in this setting. With the unimodality assumption, this led to a cutoff $\pi_{\text{thr}} = 0.87$. With the r-concavity assumption, the error bound was $\pi_{\text{thr}} = 0.69$, while the error bound became $\pi_{\text{thr}} = 1$ without assumptions. Subsequently we used cross-validation to obtain the optimal stopping iteration for the model. The code for model fitting and stability selection is given as an electronic supplement.

5.1. Results

The resulting stability paths can be found in Figure 3. The maximum inclusion frequencies for all selected base-learners and for the top scoring base-learners can be found in Figure 4. Tyrosine (`Tyr`), tryptophan (`Trp`), leucine (`Leu`) and arginine (`Arg`) all had a selection frequency of 100 %. Valine (`Val`) was selected in 97 % of the models. Without assumptions, only the amino acids with 100 % selection frequency were considered to be stable. Under the unimodality assumption, valine was additionally termed stable. Together with the sharp decline in the selection frequency, we would thus focus on these first five amino acids.

The results of our analysis using stability selection confirmed the abnormal metabolism of the amino acid tryptophan in ASD cells reported by Boccuto *et al.* (2013). Additionally, the utilization of other amino acids seemed to be affected, although on a milder level. When weighted for the size of the effect, we noticed in ASD patients an overall decreased utilization of tryptophan (-0.273 units on the logarithmic scale), tyrosine (-0.135), and valine (-0.054). On the other hand, we registered an increased rate for the metabolic utilization of arginine ($+0.084$) and leucine ($+0.081$). These findings suggest an abnormal metabolism of large amino acids (tryptophan, tyrosine, leucine, and valine), which might be related to impaired transport of those molecules across the cellular membrane. Separately, a screening by Sanger sequencing was performed on the coding regions of *SLC3A2*, *SLC7A5*, and *SLC7A8*, the genes coding the subunits of the Large Amino acid Transporter (LAT)

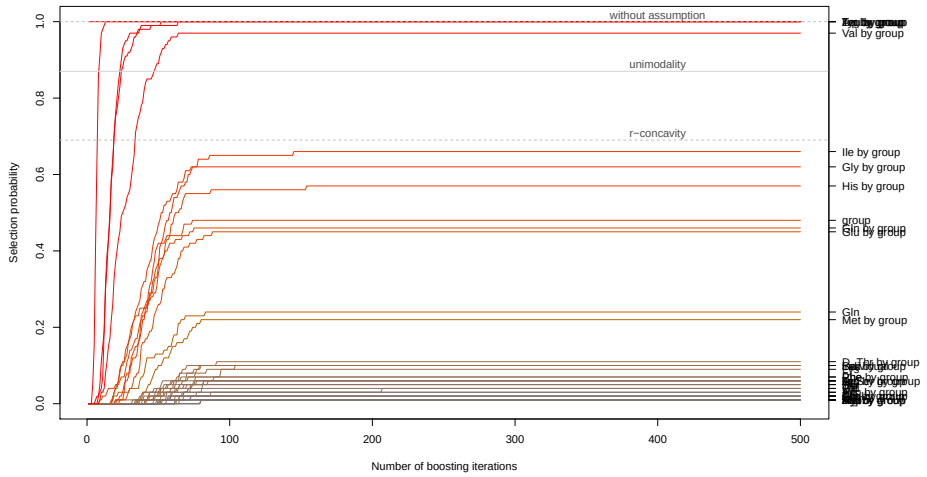

Figure 3: Stability selection paths, with the number of boosting iterations plotted against the relative selection frequency of the base-learners up to that iteration. One can deduce that the number of iterations was sufficiently large, as all selection paths cease to increase after approx. 150 iterations. The solid horizontal gray line is the threshold value with unimodality assumption ($\pi_{\text{thr}} = 0.87$), the dashed gray lines represent the threshold values with r-concavity assumption ($\pi_{\text{thr}} = 0.69$) and without assumption ($\pi_{\text{thr}} = 1$).

1 and 2, in 107 ASD patients (including the ones reported in this paper; Boccuto, unpublished data). Overall, potentially pathogenic mutations were detected in 17/107 ASD patients (15.9%): eight in *SLC3A2*, four in *SLC7A5*, and five in *SLC7A8*. We also evaluated the transcript level for these genes by expression microarray in 10 of the 17 ASD patients reported in this paper and 10 controls. The results showed that all the ASD patients had a significantly lower expression of *SLC7A5* (p value $= 0.00627$) and *SLC7A8* (p value $= 0.04067$). Therefore, we noticed that 27/107 ASD patients (25.2%) had either variants that might affect the LATs function or reduce the level of transcripts for the transporters' subunits. When we correlated the metabolic data collected by the Phenotype Microarrays with those findings, we noticed that all of these patients showed reduced utilization of tryptophan. Additionally, eight out of the twelve patients who were screened with the whole metabolic panel showed significantly reduced tyrosine utilization in at least 25 of the 27 wells containing this amino acid, seven had a reduced utilization of valine in at least 29/34 wells, and five had a reduced metabolism of leucine in at least 27/31 wells. These data are concordant with the present findings as they suggest an overall problem with the metabolism of large amino acids, which might have important consequences in neurodevelopment and synapsis homeostasis, especially if one considers that such amino acids are precursors of important compounds, such as serotonin, melatonin, quinolinic acid, and kynurenic acid (tryptophan), or dopamine (tyrosine).

6. Discussion

Stability selection proofs to work well in high-dimensional settings with (much) more predictors than observations. It adds an error control to the selection process of boosting or other high-dimensional variable selection approaches. Assumptions on the distribution of the simultaneous selection probabilities increase the number of true positive variables, while keeping the error control in most settings. As shown in our case study, complex log-linear interaction models can be used as learners in conjunction with stability selection. Additionally, more complex models such

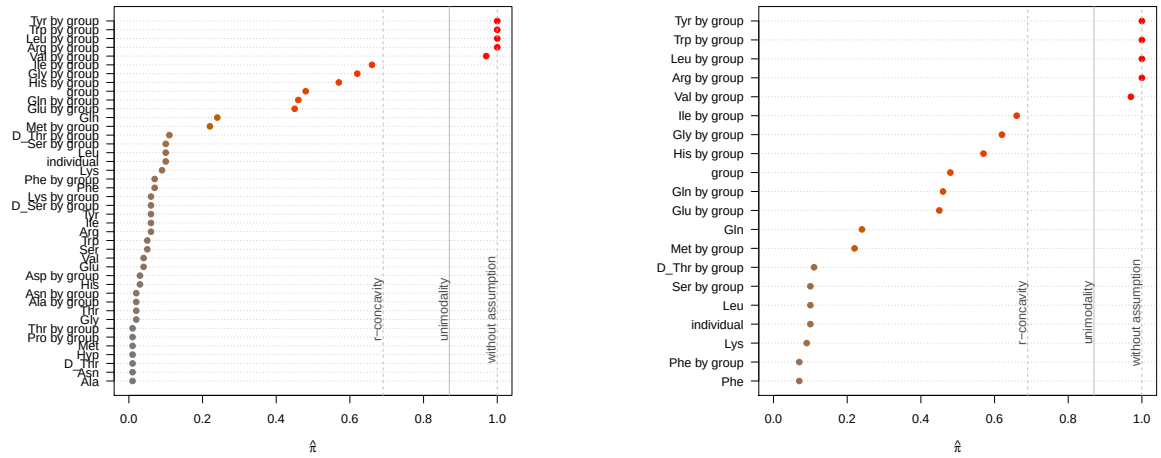

Figure 4: The maximum selection frequency $\hat{\pi}$ for all (selected) base-learners (left) and for the top 20 base-learners (right) as determined by stability selection. The solid vertical gray lines depict the threshold value with unimodality assumption ($\pi_{\text{thr}} = 0.87$), the dashed gray lines represent the threshold values with r-concavity assumption ($\pi_{\text{thr}} = 0.69$) and without assumption ($\pi_{\text{thr}} = 1$).

as generalized additive models (GAMs; Hastie and Tibshirani 1986, 1990) or structured additive regression (STAR) models Fahrmeir, Kneib, and Lang (2004); Hofner, Kneib, and Hothorn (2014a) can also benefit from the combination with stability selection if model or variable selection (with a control for the number of false positives) is of major interest.

However, one should keep in mind that stability selection controls the per-family error rate, which is very conservative. Specifying the error rate such that $\alpha \leq PFER_{\text{max}} \leq m\alpha$, with significance level α and m hypothesis tests, might provide a good idea for a sensible error control in high-dimensional settings with *FWER*-control ($PFER_{\text{max}} = \alpha$) and no multiplicity adjustment ($PFER_{\text{max}} = m\alpha$) as the extreme cases.

Furthermore, prediction models might not always benefit from stability selection. If the error control is tight, i.e. $PFER_{\text{max}}$ is small, the true positive rate is usually smaller than in a cross-validated prediction model without stability selection and the prediction accuracy suffers (see also Hothorn 2010). Prediction and variable selection are two different goals.

Implementation

The component-wise, model-based boosting approach is implemented in the R add-on package **mboost** (Bühlmann and Hothorn 2007; Hothorn *et al.* 2010; Hothorn, Bühlmann, Kneib, Schmid, and Hofner 2014). A comprehensive tutorial for **mboost** is given in Hofner *et al.* (2014b). The R package **opm** (Vaas *et al.* 2013, Göker 2014) is used to store, manage and annotate the data set. Tutorials are given as vignettes.

Stability selection is implemented in the add-on package **stabs** (Hofner and Hothorn 2014) for the statistical program environment R (R Development Core Team 2014). One can directly use stability selection on a fitted boosting model using the function `stabsel`. One only needs to additionally specify two of the parameters `PFER`, `cutoff` and `q`. The missing parameter is then computed such that the specified type of error bound holds (without additional assumptions (E1), under unimodality (E2) or under r-concavity (E3)). Alternative `stabsel` methods exist for various other fitting approaches (e.g. Lasso) using a matrix or a formula interface. By specifying a function that returns

the index (and names) of selected variables one can easily extend this framework. In general, the function `stabsel_parameters` can be used to compute the missing parameter without running stability selection itself to check if the value of the parameter computed from the other two parameters is sensible in the data situation at hand.

Acknowledgments

We thank Rajen D. Shah, N. Meinshausen and P. Bühlmann for helpful comments and discussion, Chin-Fu Chen and Charles E. Schwartz from the Greenwood Genetic Center for their help with the analysis and with the interpretation of the results, as well as Michael Drey who conducted an early version of the presented simulation study.

References

Austin PC (2008). "Bootstrap model selection had similar performance for selecting authentic and noise variables compared to backward variable elimination: a simulation study." *Journal of clinical epidemiology*, **61**, 1009–1017.

Austin PC, Tu JV (2004). "Automated variable selection methods for logistic regression produced unstable models for predicting acute myocardial infarction mortality." *Journal of clinical epidemiology*, **57**, 1138–1146.

Benjamini Y, Hochberg Y (1995). "Controlling the false discovery rate: A practical and powerful approach to multiple testing." *Journal of the Royal Statistical Society. Series B (Methodological)*, **57**, 289–300.

Boccuto L, Chen CF, Pittman A, Skinner C, McCartney H, Jones K, Bochner B, Stevenson R, Schwartz C (2013). "Decreased tryptophan metabolism in patients with autism spectrum disorders." *Molecular Autism*, **4**(1), 16.

Bochner BR, Gadzinski P, Panomitros E (2001). "Phenotype microarrays for high throughput phenotypic testing and assay of gene function." *Genome Research*, **11**, 1246–1255.

Breiman L (2001). "Random forests." *Machine Learning*, **45**, 5–32.

Bühlmann P, Hothorn T (2007). "Boosting algorithms: Regularization, prediction and model fitting." *Statistical Science*, **22**, 477–505.

Bühlmann P, Yu B (2003). "Boosting with the L_2 loss: Regression and classification." *Journal of the American Statistical Association*, **98**, 324–339.

Bühlmann P, Kalisch M, Meier L (2014). "High-dimensional statistics with a view toward applications in biology." *Annual Review of Statistics and Its Application*, **1**, 255–278.

Chaturvedi N, Goeman J, Boer J, van Wieringen W, de Menezes R (2014). "A test for comparing two groups of samples when analyzing multiple omics profiles." *BMC Bioinformatics*, **15**(1), 236.

Dudoit S, Shaffer JP, Boldrick JC (2003). "Multiple hypothesis testing in microarray experiments." *Statistical Science*, **18**, 71–103.

Efron B, Hastie T, Johnstone I, Tibshirani R (2004). "Least angle regression (with discussion)." *The Annals of Statistics*, **32**, 407–451.

Fahrmeir L, Kneib T, Lang S (2004). "Penalized structured additive regression: A Bayesian perspective." *Statistica Sinica*, **14**, 731–761.

Fellinghauer B, Bühlmann P, Ryffel M, von Rhein M, Reinhardt JD (2013). "Stable graphical model estimation with random forests for discrete, continuous, and mixed variables." *Computational Statistics & Data Analysis*, **64**, 132 – 152.

Flack VF, Chang PC (1987). "Frequency of selecting noise variables in subset regression analysis: a simulation study." *The American Statistician*, **41**, 84–86.

Friedman J, Hastie T, Tibshirani R (2000). "Additive logistic regression: A statistical view of boosting (with discussion)." *The Annals of Statistics*, **28**, 337–407.

Göker M, with contributions by B Hofner, LAI Vaas, J Sikorski, N Buddruhs and A Fiebig (2014). *opm: Analysing phenotype microarray and growth curve data*. R package version 1.1-0, URL `http://CRAN.R-project.org/package=opm`.

Groth P, Weiss B, Pohlenz HD, Leser U (2008). "Mining phenotypes for gene function prediction." *BMC Bioinformatics*, **9**(1), 136.

Hastie T, Tibshirani R (1986). "Generalized additive models." *Statistical Science*, **1**, 297–310.

Hastie T, Tibshirani R (1990). *Generalized additive models*. Chapman & Hall / CRC, London.

Haury AC, Mordelet F, Vera-Licona P, Vert JP (2012). "TIGRESS: Trustful Inference of Gene REgulation using Stability Selection." *BMC Systems Biology*, **6**(1), 145.

He Q, Lin DY (2011). "A variable selection method for genome-wide association studies." *Bioinformatics*, **27**(1), 1–8.

Hofner B, Hothorn T (2014). *stabs: Stability selection with error control*. R package version 0.1-0, URL `http://CRAN.R-project.org/package=stabs`.

Hofner B, Hothorn T, Kneib T, Schmid M (2011). "A framework for unbiased model selection based on boosting." *Journal of Computational and Graphical Statistics*, **20**, 956–971.

Hofner B, Kneib T, Hothorn T (2014a). "A unified framework of constrained regression." *Statistics and Computing*, pp. 1–14. doi:10.1007/s11222-014-9520-y.

Hofner B, Mayr A, Robinzonov N, Schmid M (2014b). "Model-based boosting in R – A hands-on tutorial using the R package mboost." *Computational Statistics*, **29**, 3–35.

Hothorn T (2010). "Discussion: Stability selection." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **72**, 463–464.

Hothorn T, Bühlmann P, Kneib T, Schmid M, Hofner B (2010). "Model-based boosting 2.0." *Journal of Machine Learning Research*, **11**, 2109–2113.

Hothorn T, Bühlmann P, Kneib T, Schmid M, Hofner B (2014). *mboost: Model-based boosting*. R package version 2.4-0, URL `http://CRAN.R-project.org/package=mboost`.

Hothorn T, Müller J, Schröder B, Kneib T, Brandl R (2011). "Decomposing environmental, spatial, and spatiotemporal components of species distributions." *Ecological Monographs*, **81**, 329–347.

Kneib T, Hothorn T, Tutz G (2009). "Variable selection and model choice in geoadditive regression models." *Biometrics*, **65**, 626–634.

Lindon JC, Holmes E, Nicholson JK (2003). "So what's the deal with metabonomics?" *Analytical Chemistry*, **75**, 385–391.

Ludwig C, Günther UL (2011). "Metabolab: Advanced NMR data processing and analysis for metabolomics." *BMC Bioinformatics*, **12**(1), 366.

Mallick P, Kuster B (2010). "Proteomics: A pragmatic perspective." *Nature Biotechnology*, **28**(7), 695–709. doi:10.1038/nbt.1658.

Manning-Courtney P, Murray D, Currans K, Johnson H, Bing N, Kroeger-Geoppinger K, Sorensen R, Bass J, Reinhold J, Johnson A, Messerschmidt T (2013). "Autism spectrum disorders." *Current Problems in Pediatric and Adolescent Health Care*, **43**(1), 2 – 11. Autism Spectrum Disorders.

Marbach D, Costello JC, Küffner R, Vega NM, Prill RJ, Camacho DM, Allison KR, Kellis M, Collins JJ, Stolovitzky G, *et al.* (2012). "Wisdom of crowds for robust gene network inference." *Nature methods*, **9**(8), 796–804.

Mayr A, Hofner B, Schmid M (2012). "The importance of knowing when to stop – a sequential stopping rule for component-wise gradient boosting." *Methods of Information in Medicine*, **51**, 178–186.

Meinshausen N, Bühlmann P (2010). "Stability selection (with discussion)." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **72**, 417–473.

R Development Core Team (2014). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0, URL `http://www.R-project.org`.

Schmid M, Hothorn T (2008). "Boosting additive models using component-wise P-splines." *Computational Statistics & Data Analysis*, **53**, 298–311.

Schmid M, Hothorn T, Krause F, Rabe C (2012). "A PAUC-based estimation technique for disease classification and biomarker selection." *Statistical Applications in Genetics and Molecular Biology*. Accepted.

Shah RD, Samworth RJ (2013). "Variable selection with error control: another look at stability selection." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **75**, 55–80.

Strobl C, Boulesteix AL, Zeileis A, Hothorn T (2007). "Bias in random forest variable importance measures: Illustrations, sources and a solution." *BMC Bioinformatics*, **8**, 25.

Tibshirani R (1996). "Regression shrinkage and selection via the lasso." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **58**, 267–288.

Vaas LAI, Sikorski J, Hofner B, Buddruhs N, Fiebig A, Klenk HP, Göker M (2013). "opm: An R package for analysing OmniLog® phenotype microarray data." *Bioinformatics*, **29**(14), 1823–1824.

Wang Z, Gerstein M, Snyder M (2009). "RNA-Seq: A revolutionary tool for transcriptomics." *Nature Reviews Genetics*, **10**(1), 57–63. doi:10.1038/nrg2484.

Zou H, Hastie T (2005). "Regularization and variable selection via the elastic net." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **67**, 301–320.

A. Additional Figures

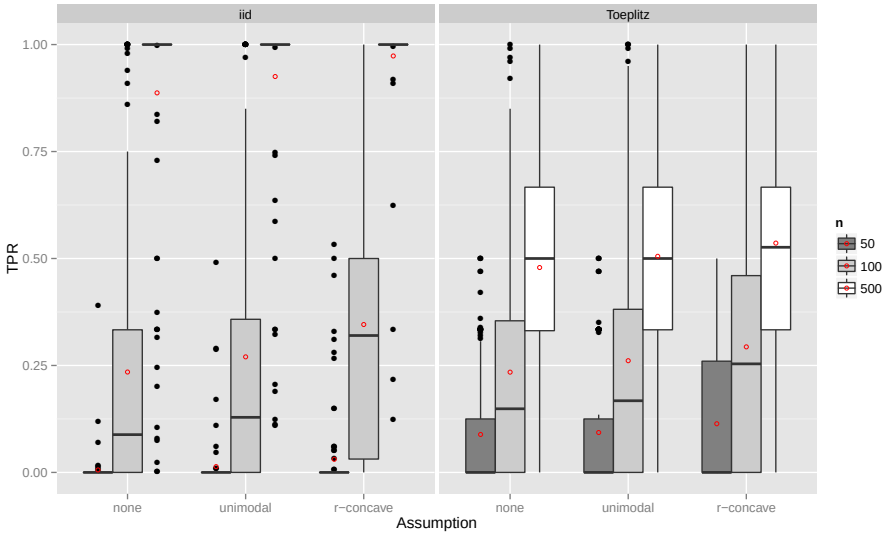

Figure 5: Boxplots for the true positives rates (TPR) for all simulation settings with separate boxplots for different numbers of observations (n), the correlation settings (independent predictor variables or Toeplitz design), and the assumptions used to compute the error bound. Each observation in the boxplot is the average of the 50 simulation replicates. The open red circles represent the average true positive rates.

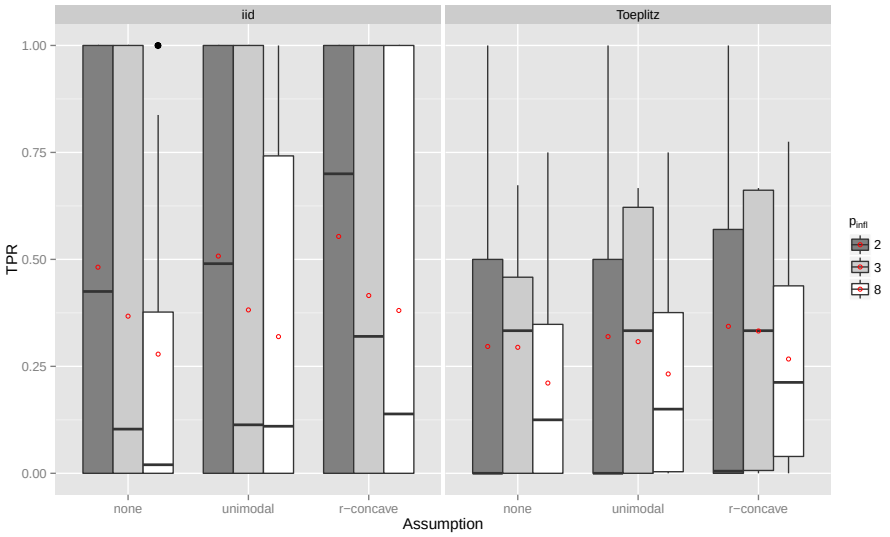

Figure 6: Boxplots for the true positives rates (TPR) for all simulation settings with separate boxplots for different numbers of influential variables (p_{infl}), the correlation settings (independent predictor variables or Toeplitz design), and the assumptions used to compute the error bound. Each observation in the boxplot is the average of the 50 simulation replicates. The open red circles represent the average true positive rates.

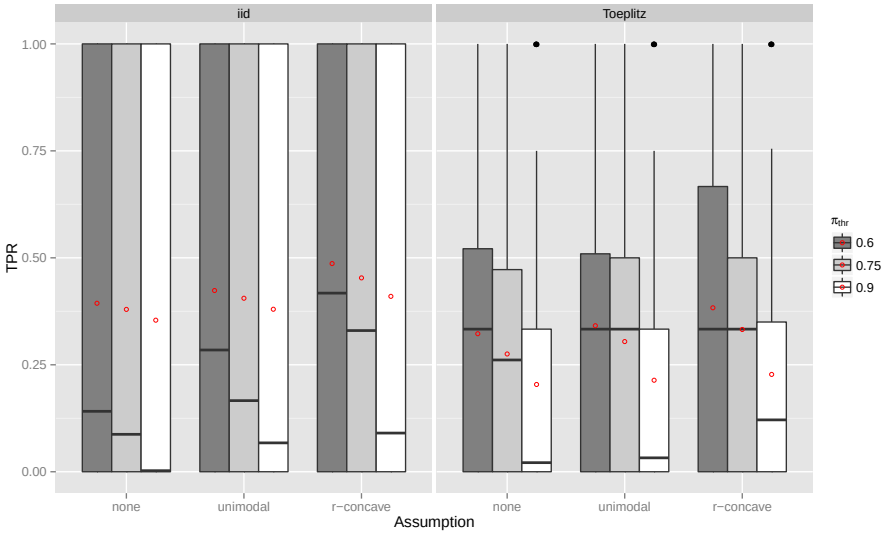

Figure 7: Boxplots for the true positives rates (TPR) for all simulation settings with separate boxplots for different cutoff values (π_{thr}), the correlation settings (independent predictor variables or Toeplitz design), and the assumptions used to compute the error bound. Each observation in the boxplot is the average of the 50 simulation replicates. The open red circles represent the average true positive rates.

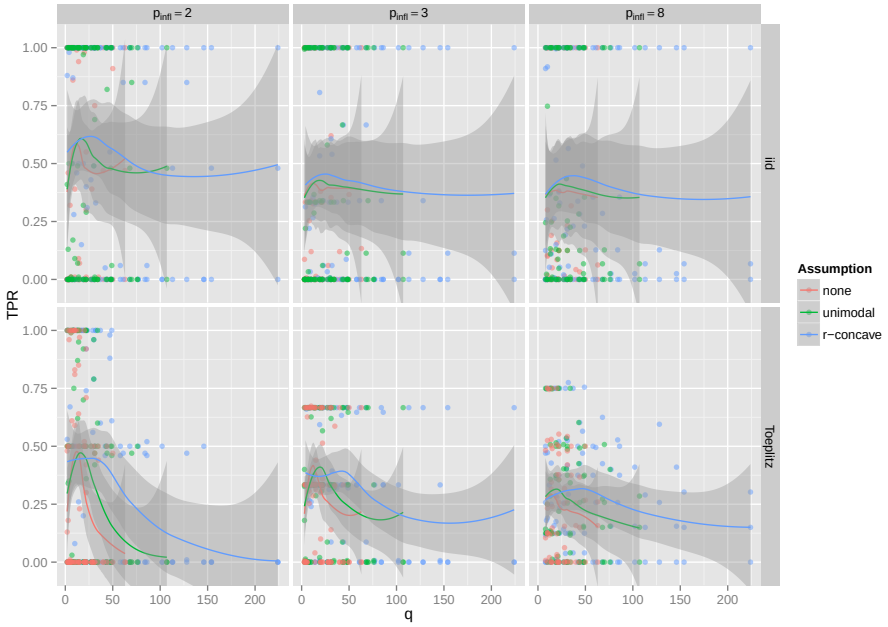

Figure 8: Scatter plot showing the true positives rates (TPR) for all simulation settings where q was larger than the number of influential variables (p_{infl}); Plots are shown separately for the number of influential variables (p_{infl}), the correlation settings (independent predictor variables or Toeplitz design) and the assumptions used to compute the error bound. Each observation in the plot is the average of the 50 simulation replicates. The lines depict a scatter plot smoother for each group together with shaded confidence regions.

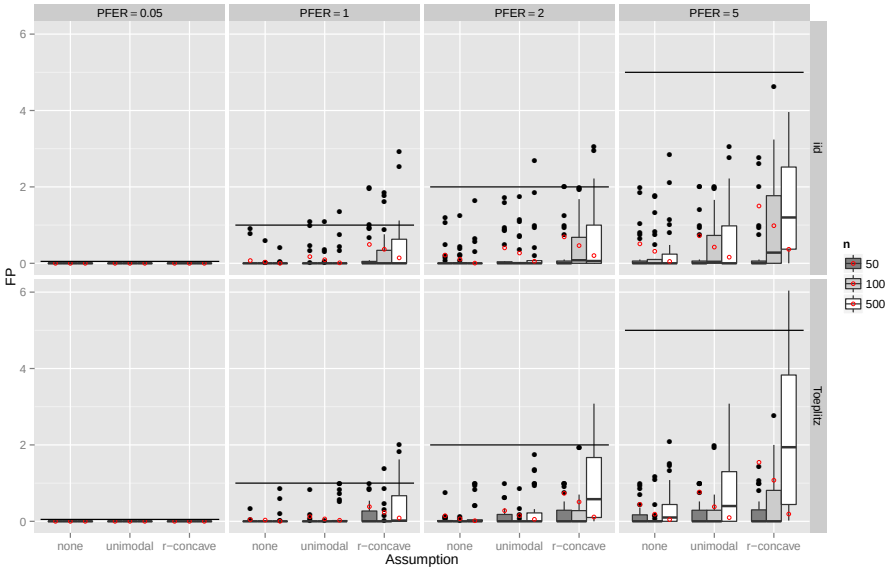

Figure 9: Boxplots for the number of false positives (FP) for all simulation settings with separate boxplots for different numbers of observations (n), the correlation settings (independent predictor variables or Toeplitz design), the *PFER*, and the assumptions used to compute the error bound. Each observation in the boxplot is the average of the 50 simulation replicates. The open red circles represent the average number of false positives.

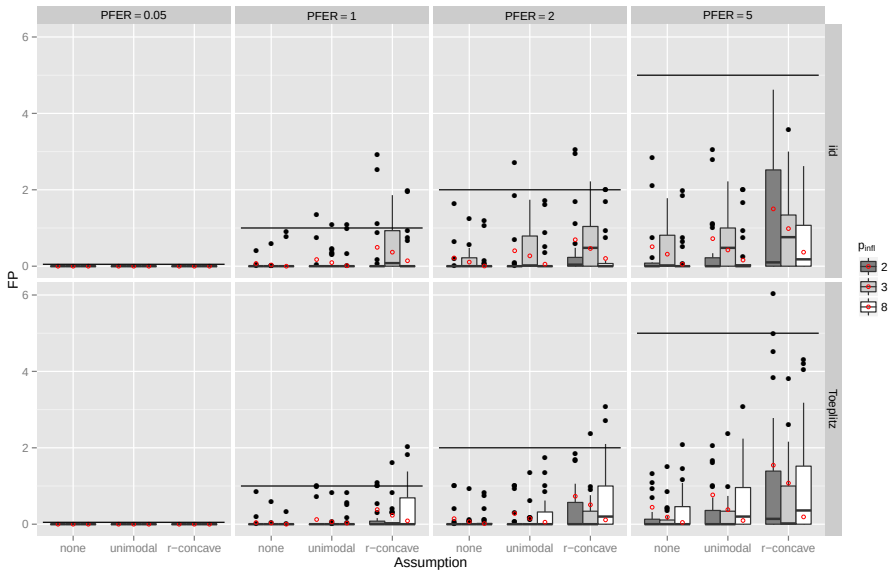

Figure 10: Boxplots for the number of false positives (FP) for all simulation settings with separate boxplots for different numbers of influential variables (p_{infl}), the correlation settings (independent predictor variables or Toeplitz design), the *PFER*, and the assumptions used to compute the error bound. Each observation in the boxplot is the average of the 50 simulation replicates. The open red circles represent the average number of false positives.

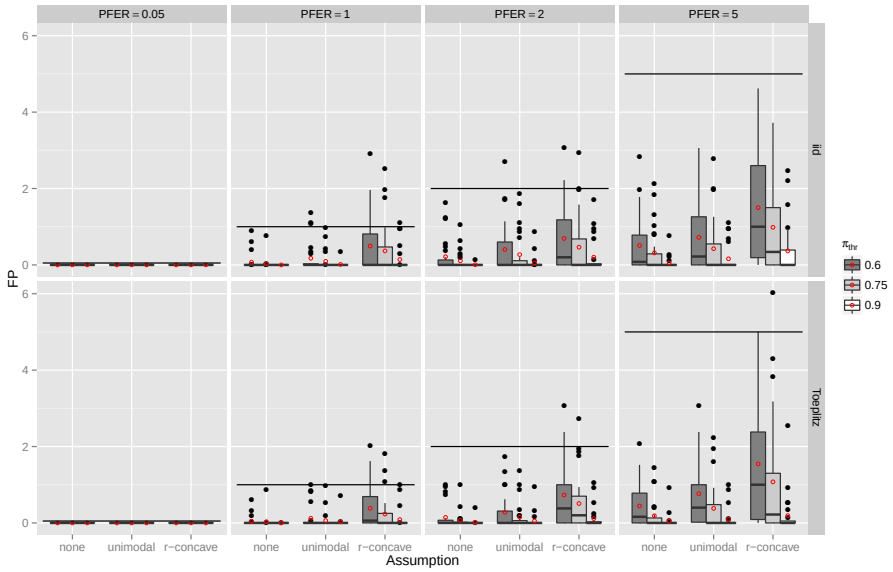

Figure 11: Boxplots for the number of false positives (FP) for all simulation settings with separate boxplots for different cutoff values (π_{thr}), the correlation settings (independent predictor variables or Toeplitz design), the *PFER*, and the assumptions used to compute the error bound. Each observation in the boxplot is the average of the 50 simulation replicates. The open red circles represent the average number of false positives.

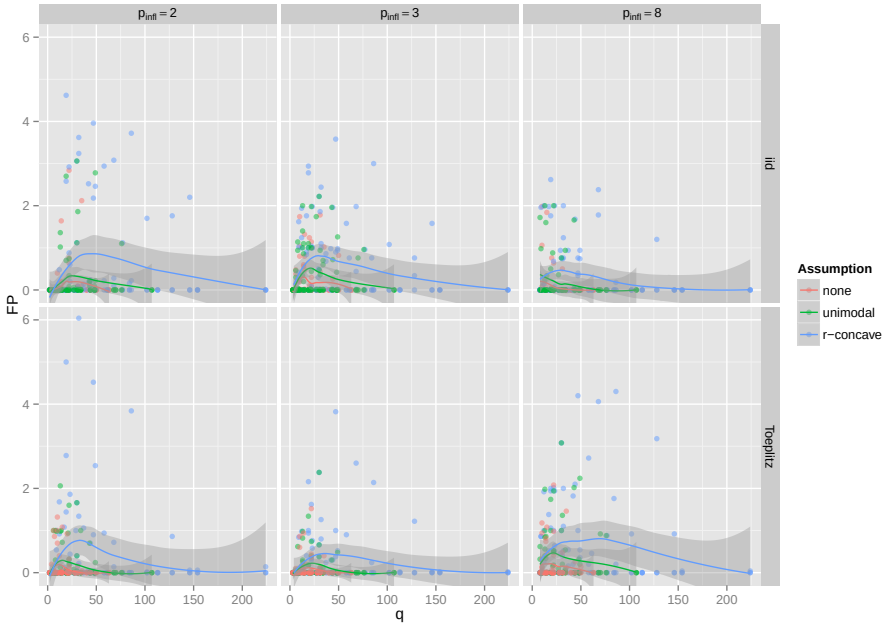

Figure 12: Scatter plot showing the number of false positives (FP) for all simulation settings where q was larger than the number of influential variables (p_{infl}); Plots are shown separately for the number of influential variables (p_{infl}), the correlation settings (independent predictor variables or Toeplitz design) and the assumptions used to compute the error bound. Each observation in the plot is the average of the 50 simulation replicates. The lines depict a scatter plot smoother for each group together with shaded confidence regions.